



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

A Collaboration and Commercialization Model for Exascale Software Research

M. Seager, B. Gorda

July 22, 2009

International Journal of High Performance Computing
Applications

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

A Collaboration and Commercialization Model for Exascale Software Research

Mark Seager and Brent Gorda,
Lawrence Livermore National Laboratory
June 9, 2009
Version 4

Motivation

In the US, recent software research and development for petascale systems has been performed by two main entities: US Government funded R&D collaborations (both at Universities and at Government Labs) and Industry efforts at products. With few notable exceptions, there has been little diffusion of technology from the R&D collaborations to industrial efforts and little feedback from the industrial efforts to the US government funded R&D efforts. However, the broader community has found value in some of the R&D efforts and would like to see continued support. For the most part, support is voluntary by the development groups because the funding was only for the R&D, not ongoing support. On the other hand, industry efforts end up being funded for specific platforms and are generally proprietary and suffer from the lack of overall effort due limited private and public investment. Understanding these lessons from petascale efforts is essential for forming a coherent strategy going forward to exascale. Clearly, a different research and development and commercialization model is desired going forward.

Proposed Model

Many US Government funded R&D collaborations produce useful results and lessons learned that are available to the HPC community for a variety of platforms. There is also much duplication of effort within various HPC vendor organizations in the name of differentiation and specialization. Both of these approaches are inefficient because they don't effectively leverage each other. The basic R&D efforts don't feed into commercial development models and overall requirements from customers fielding systems are not being fed back into the R&D efforts.

To overcome this and align forces toward the Petascale, we propose a new Open Source Collaborative R&D model with commercialization paths. This leverages the "best of breed" development models from DOE Office of Science (DOE/SC) petascale research efforts that are typically Open Source, Community development based. It also leverages the NNSA Advanced Computing and Simulation (ASC) PathForward (now FastForward) program where HPC provider product roadmaps are accelerated and provide a clear commercialization strategy.

Figure 1 depicts the proposed model graphically. In this model for software development for exascale systems, we retain the flexibility of R&D efforts to experiment, push the boundary and to be allowed to fail. The fruits of these efforts (in the blue STAR figure) are handed off as harvestable results (e.g., code,

algorithms, models or techniques) and as “lessons learned.” These are harvested by a new class of efforts labeled as Development and Engineering (D&E) collaborations in the Orange Box. These D&E ASC PathForward-like efforts include a commercialization path should the results be successful. These products are then delivered and supported on various HPC systems by the providers of these commercial technologies (e.g., system software by system vendors and ISV products such as code development tools). The key difference is the management and funding model for these efforts. Rather than separate independent efforts in R&D, D&E, Products and Support, we propose they be linked. Funding agencies for the D&E collaborations (E.g., ASC and DARPA) should participate as contributors in the R&D efforts (e.g., DOE SC and NSF). That is, the R&D organizations should continue to lead the R&D portions, but include contributions from organizations that focus on the D&E collaborations. Likewise, R&D organizations should contribute to the D&E funding planning and execution in the D&E efforts. As vendor partners contribute to the D&E collaborations, natural commercialization strategies will emerge. Vendor partners should also be included, when appropriate, in the R&D collaborations.

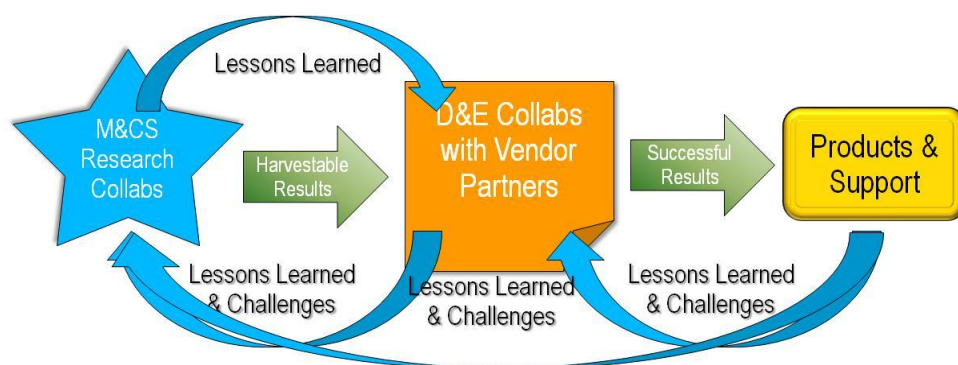


Figure 1: A new software development model for exascale systems couples basic R&D with commercial effort so leverage the best of both worlds.

In all cases, linkages between stages should be valued as part of the project selection process in order to incentivize the migration of technology from R&D to D&E and ultimately into products and services. Naturally some R&D proposals could be formed without D&E collaboration paths, but may be selected for funding based on the strength of the technical merits. In other words, the model should be flexible, but encourage and incentivize technology migration.

A side effect of this strategy is that at every stage of migrating technology from left to right in Figure 1, there is a corresponding opportunity to shape the agenda of upstream events by migrating challenges, requirements and “Lessons learned” in counter flow direction (right to left in Figure 1).

There is a large gap between what has been developed for current 100s of teraFLOP/s Linux clusters and 1-20 petaFLOP/s systems that have been delivered or are on the horizon. The larger system comes with huge requirements in terms of scalable systems software and file systems; Reliability Availability and Serviceability (RAS); programming models and application resiliency. It is important that the community consider multiple passes through the process depicted in Figure 1 be attempted before fielding exascale systems in 2018 and beyond.

MAIN PRINCIPLES

1. Coordinate strategy between R&D->D&E and D&E->P&S. With migration path towards commercialization.
2. Keep current focus areas and funding agents for R&D, D&E and P&S as they currently are and add stake holders from next stage in the process.
3. Keep the model flexible as possible to encourage development and competition.
4. Multiple iterations required to get to exascale.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.